Itanium® 2 Processor Microarchitecture Overview

Terry Lyon

presented
June 24, 2002
at the
Vail Computer Elements Workshop

Itanium® 2 Processor Microarchitecture

Block Diagram

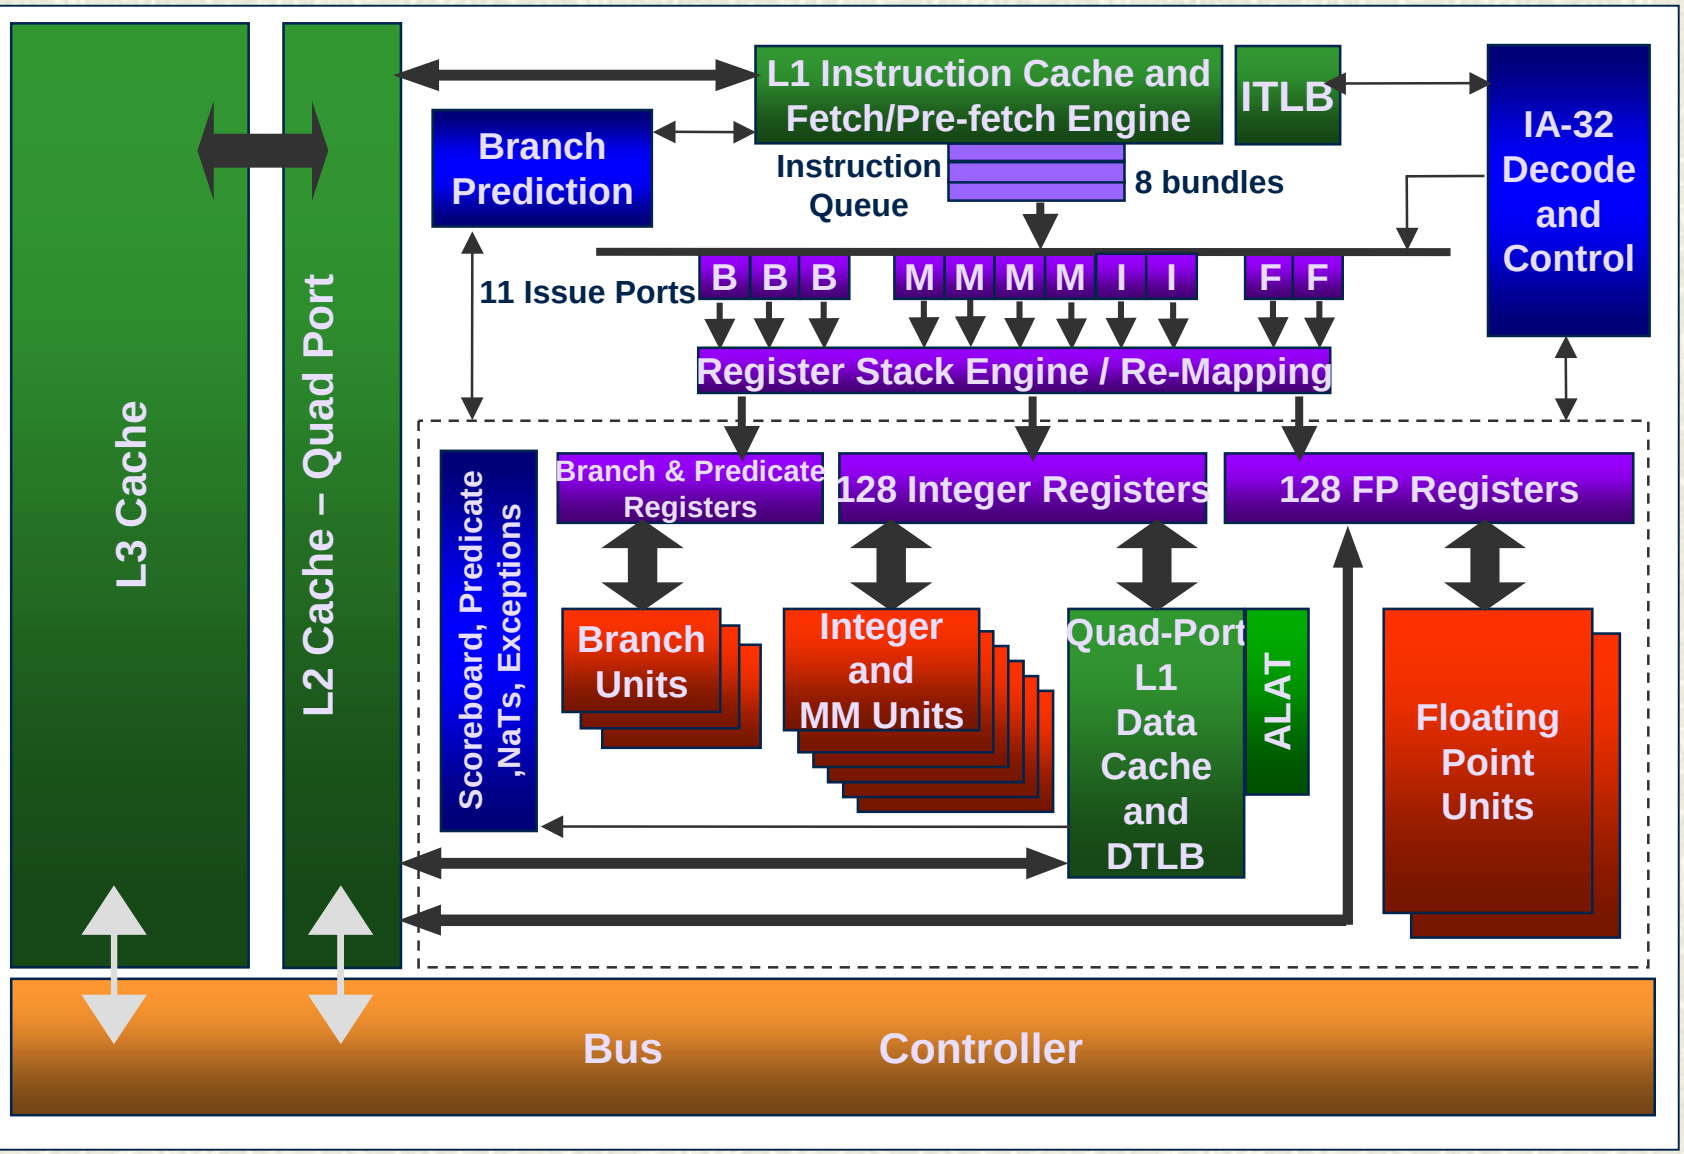

Itanium® 2 Processor Microarchitecture

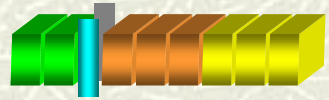

Main Pipeline

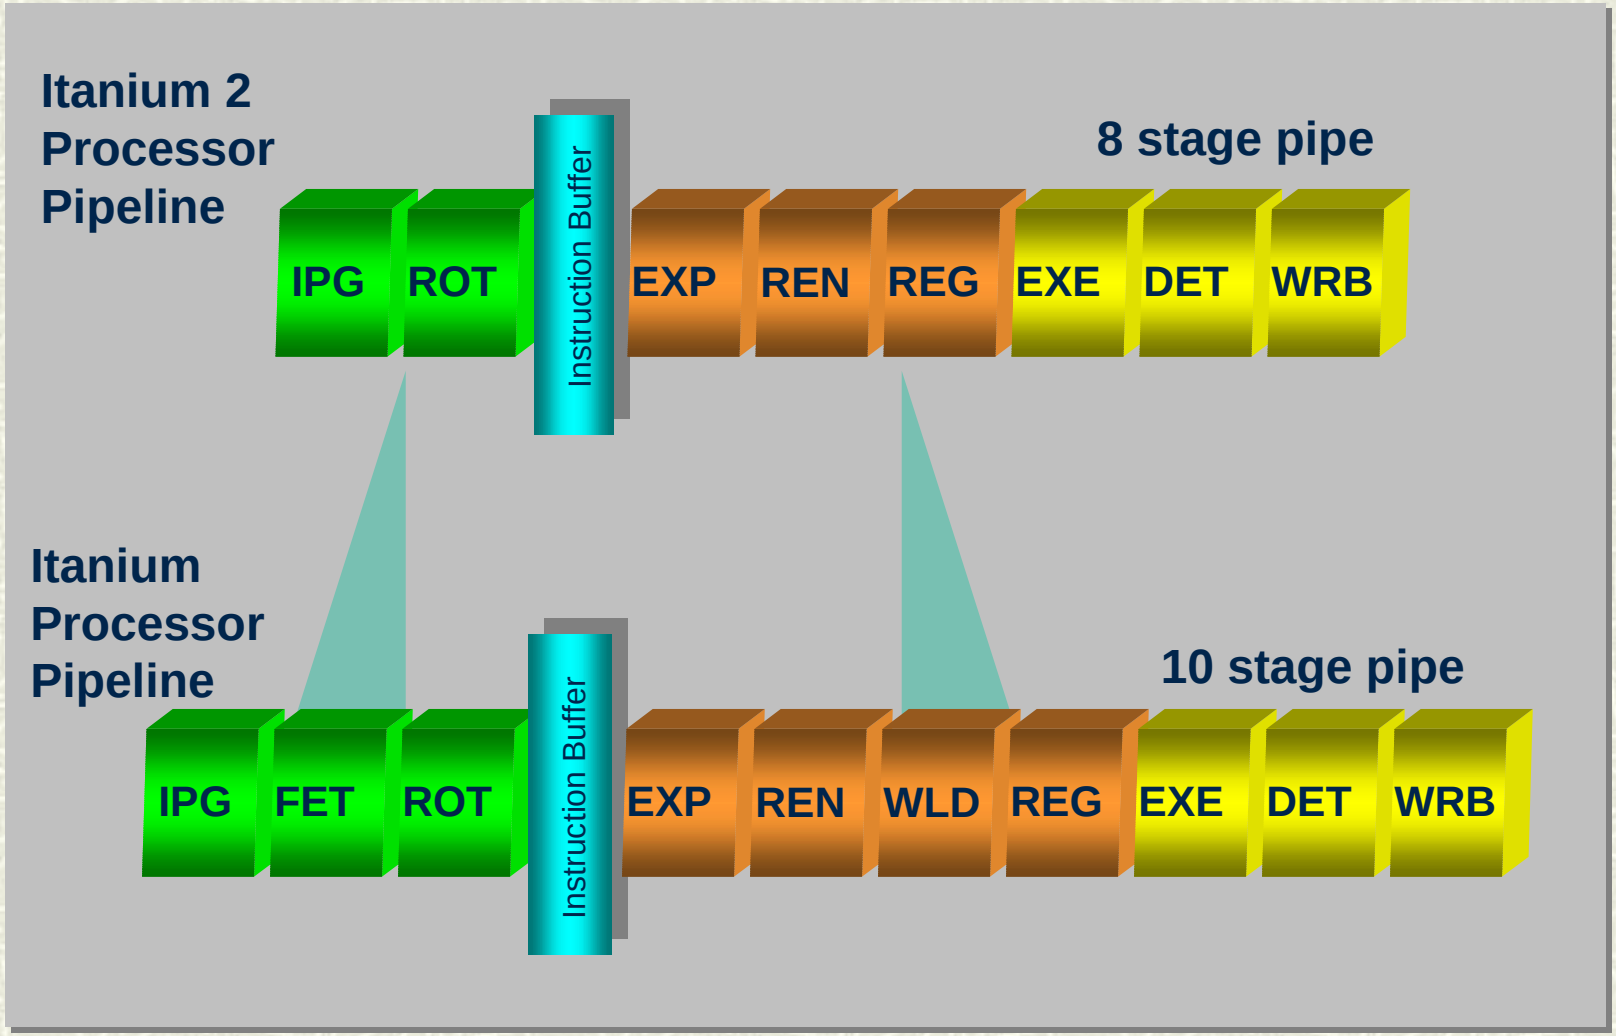

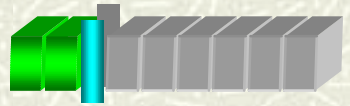

Instruction Fetch Features

- L1i Cache
  - 16KB, 4 way assoc, 64B line size
  - 1 cycle access (reduces pipe stage)
  - prevalidated tag design
- Fill bandwidth from L2/L3 equals 2 bundles/cycle
- 8 bundle **Instruction Buffer** to smooth flow to main pipeline

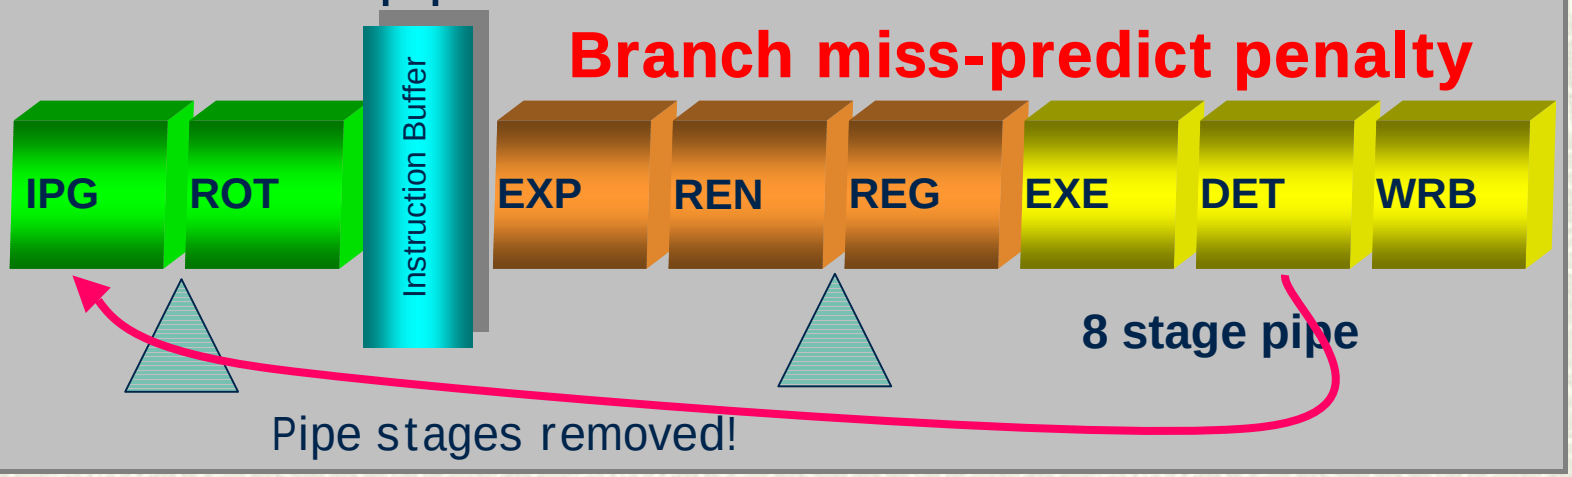

Itanium® 2 Processor Microarchitecture

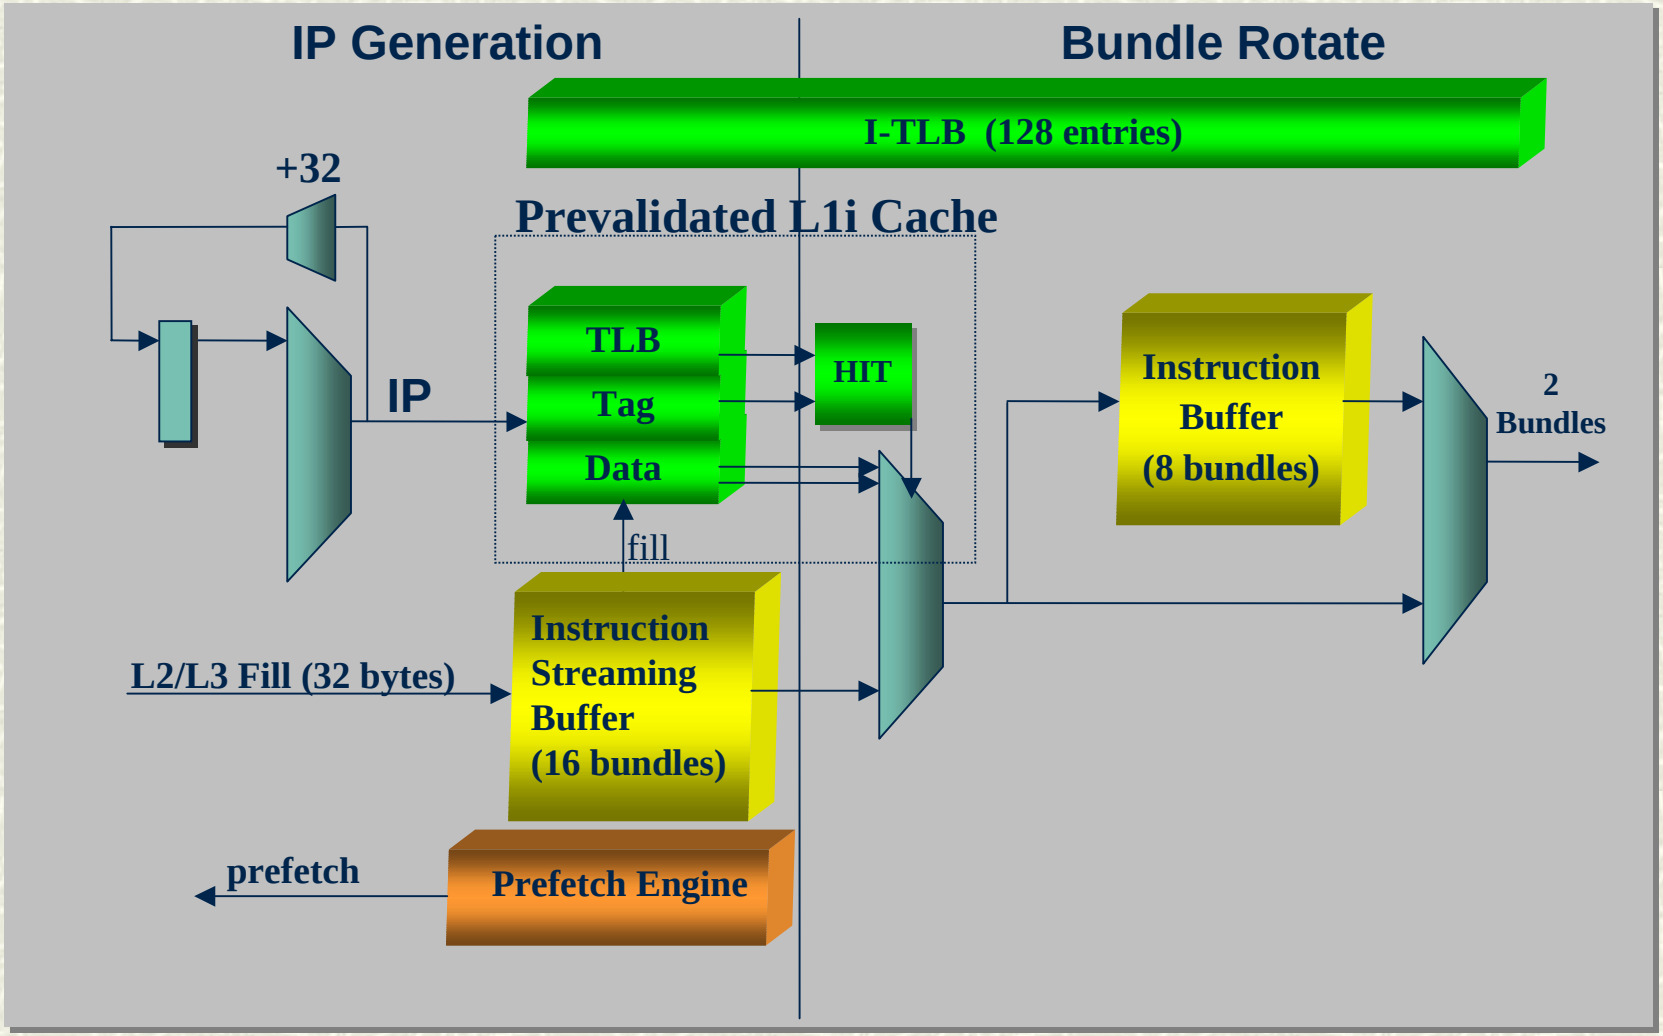

I-Cache / I-Fetch

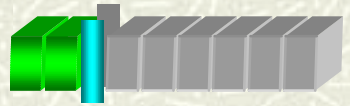

Instruction prefetching

- Software controls prefetch hardware
- Separate demand and prefetch L1i tag ports allow parallel lookups
- Up to 8 I miss requests may be outstanding

Instruction	Class	Prefetch range
any branch	Demand	0-12 Instructions
brp.few	Hint	12-24 Instructions
brp.many	Hint	36-48 Instructions
mov br.few	Hint	12-24 Instructions
br.many	Stream	Stream until redirected

Fast Branch Resteer

- 0-cycle IP relative branch resteer
- branch info in L1i cache

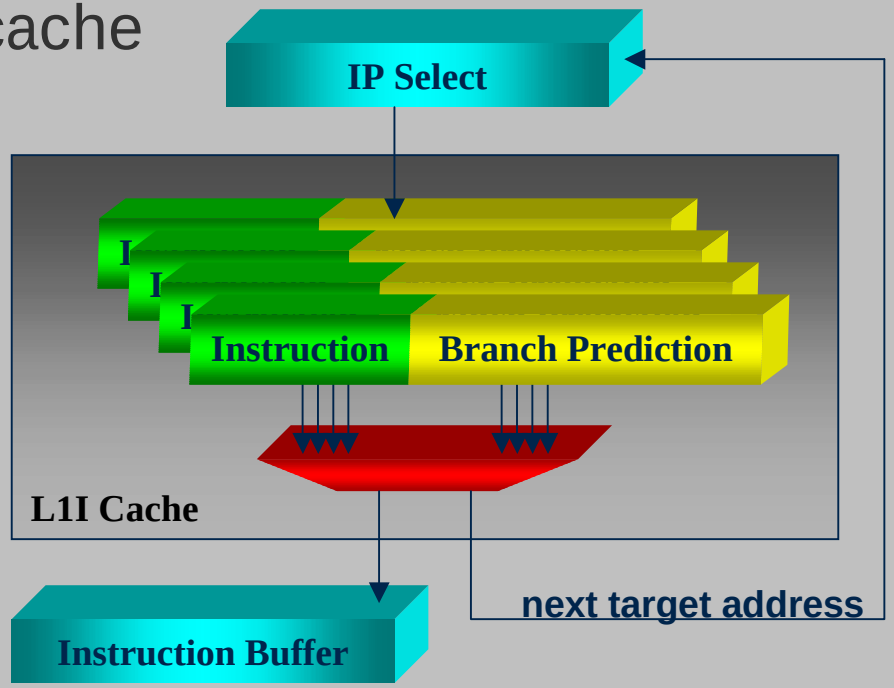

Itanium® 2 Processor Microarchitecture

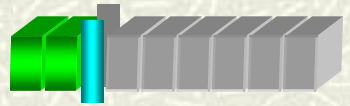

Branch Prediction Structures

- for IP relative br →L1i and Instr. Steaming Buffer
- for returns → Return Stack Buffer (8 entry)
- for indirect branches →
  Speculative Branch Register (8 entry)
- for loop branches → Perfect Loop Prediction

Supporting Structures

- Pattern History Table (predicts for L1/ISB)
  - generates next branch prediction for the L1 cache entries
  - holds 16K 2bit counts
- L2B (holds branch information not in L1i)
  - second level branch history structure
  - used to supply prediction information for new L1i entries

Itanium® 2 Processor Microarchitecture

Dispersal

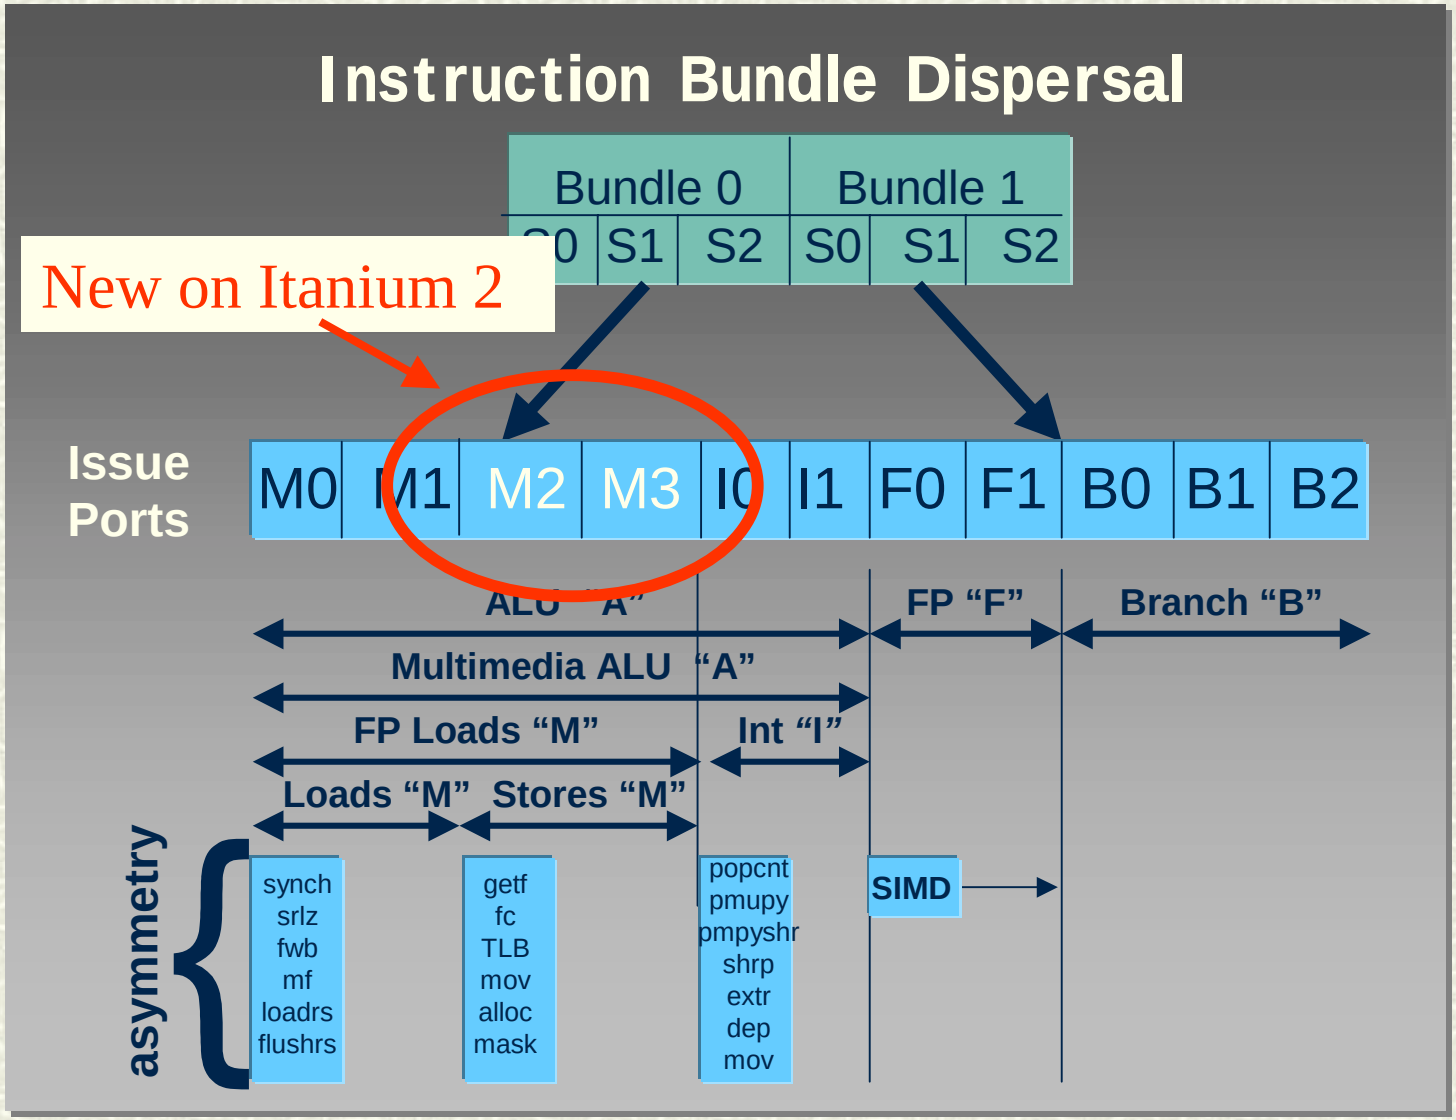

Itanium® 2 Processor Microarchitecture

Dispersal

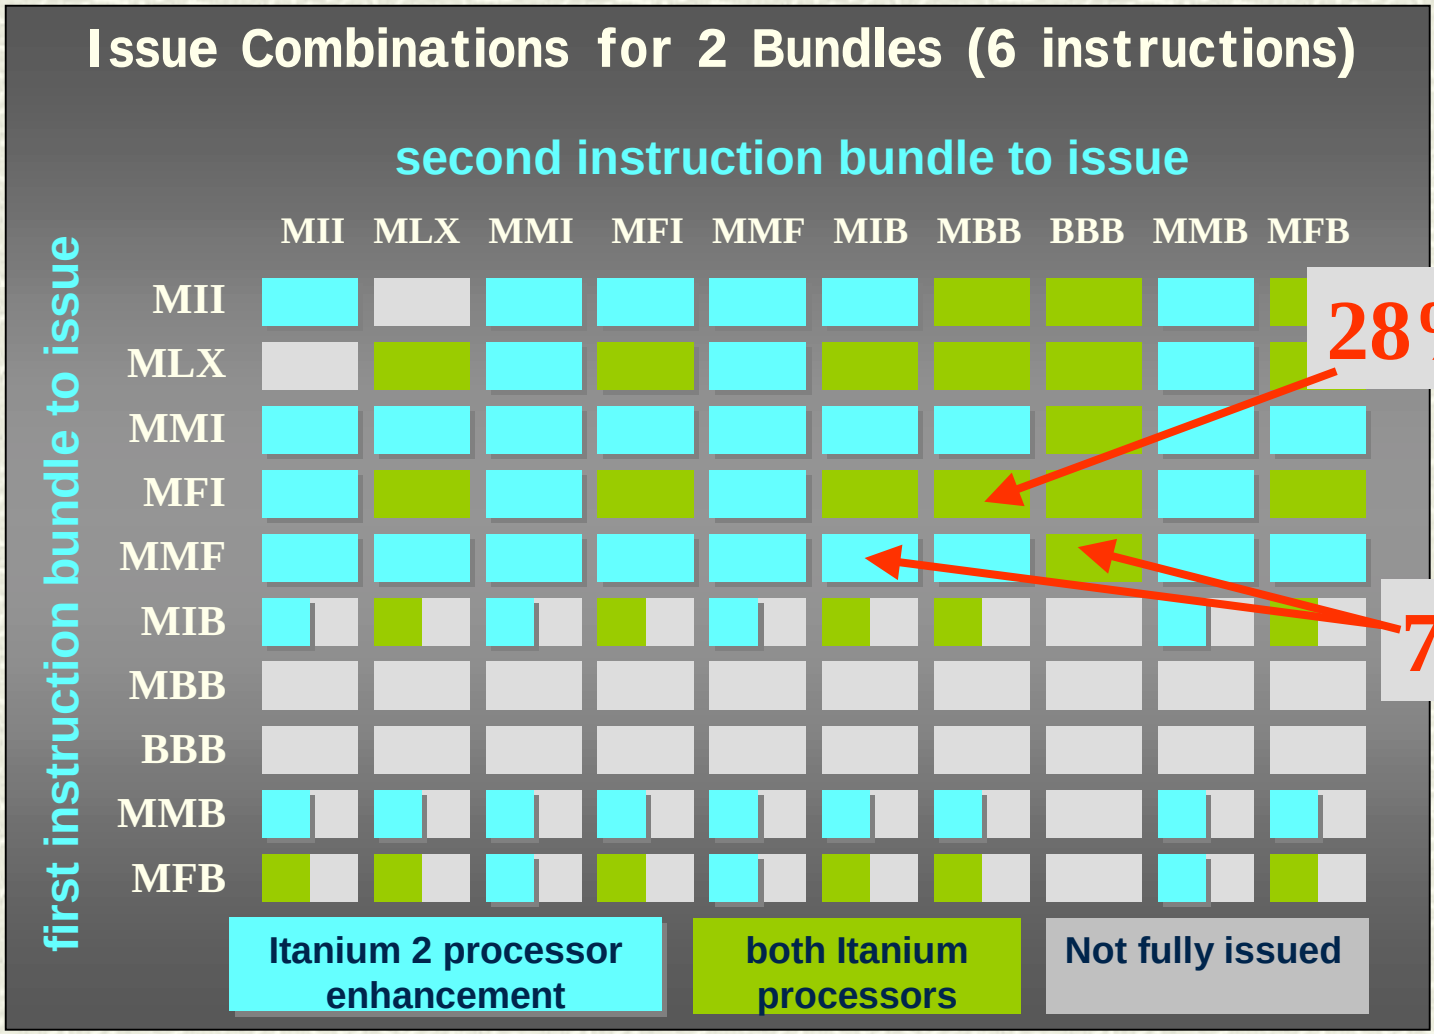

Itanium® 2 Processor Microarchitecture

General Register File

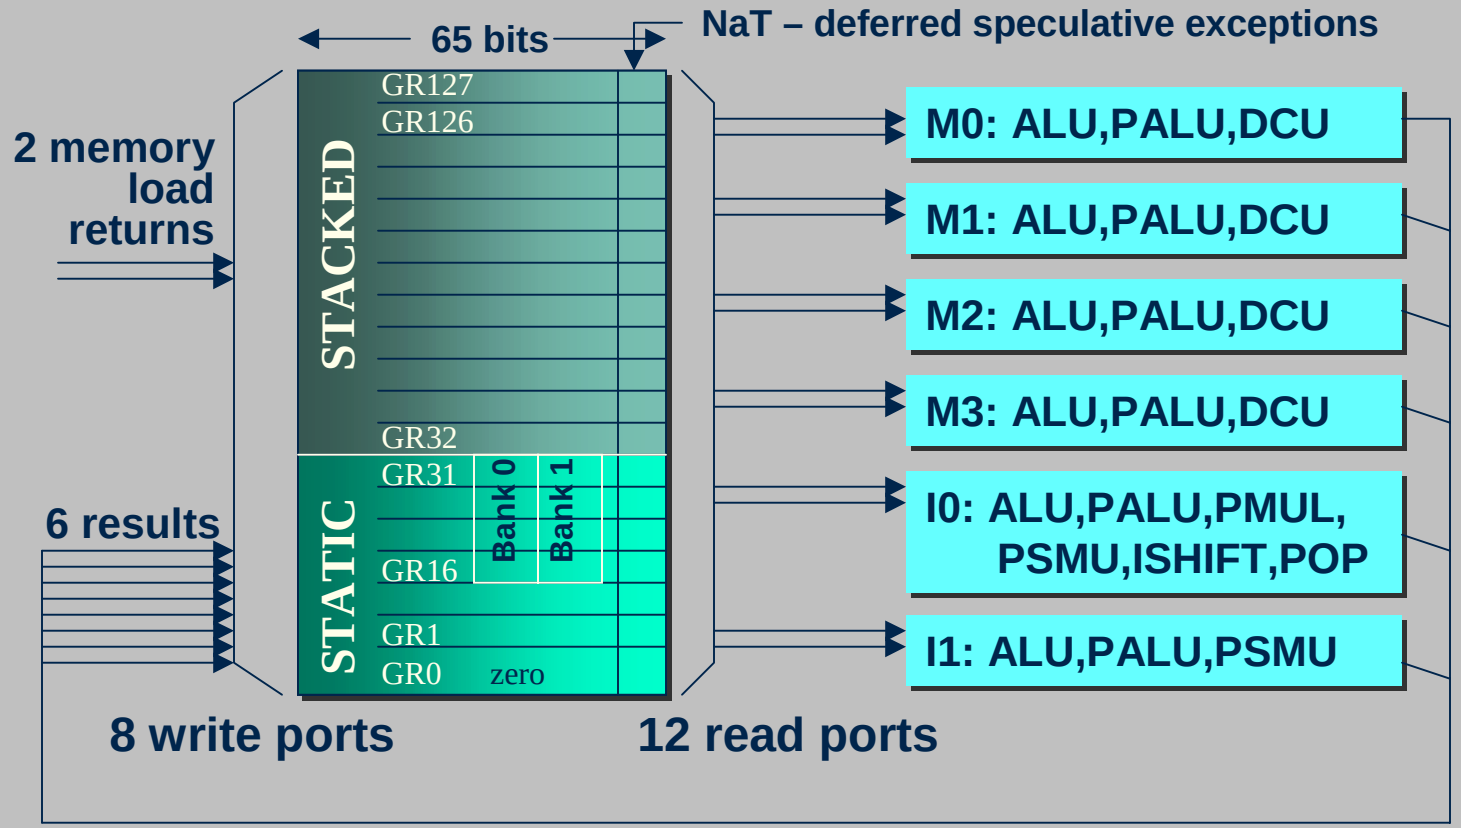

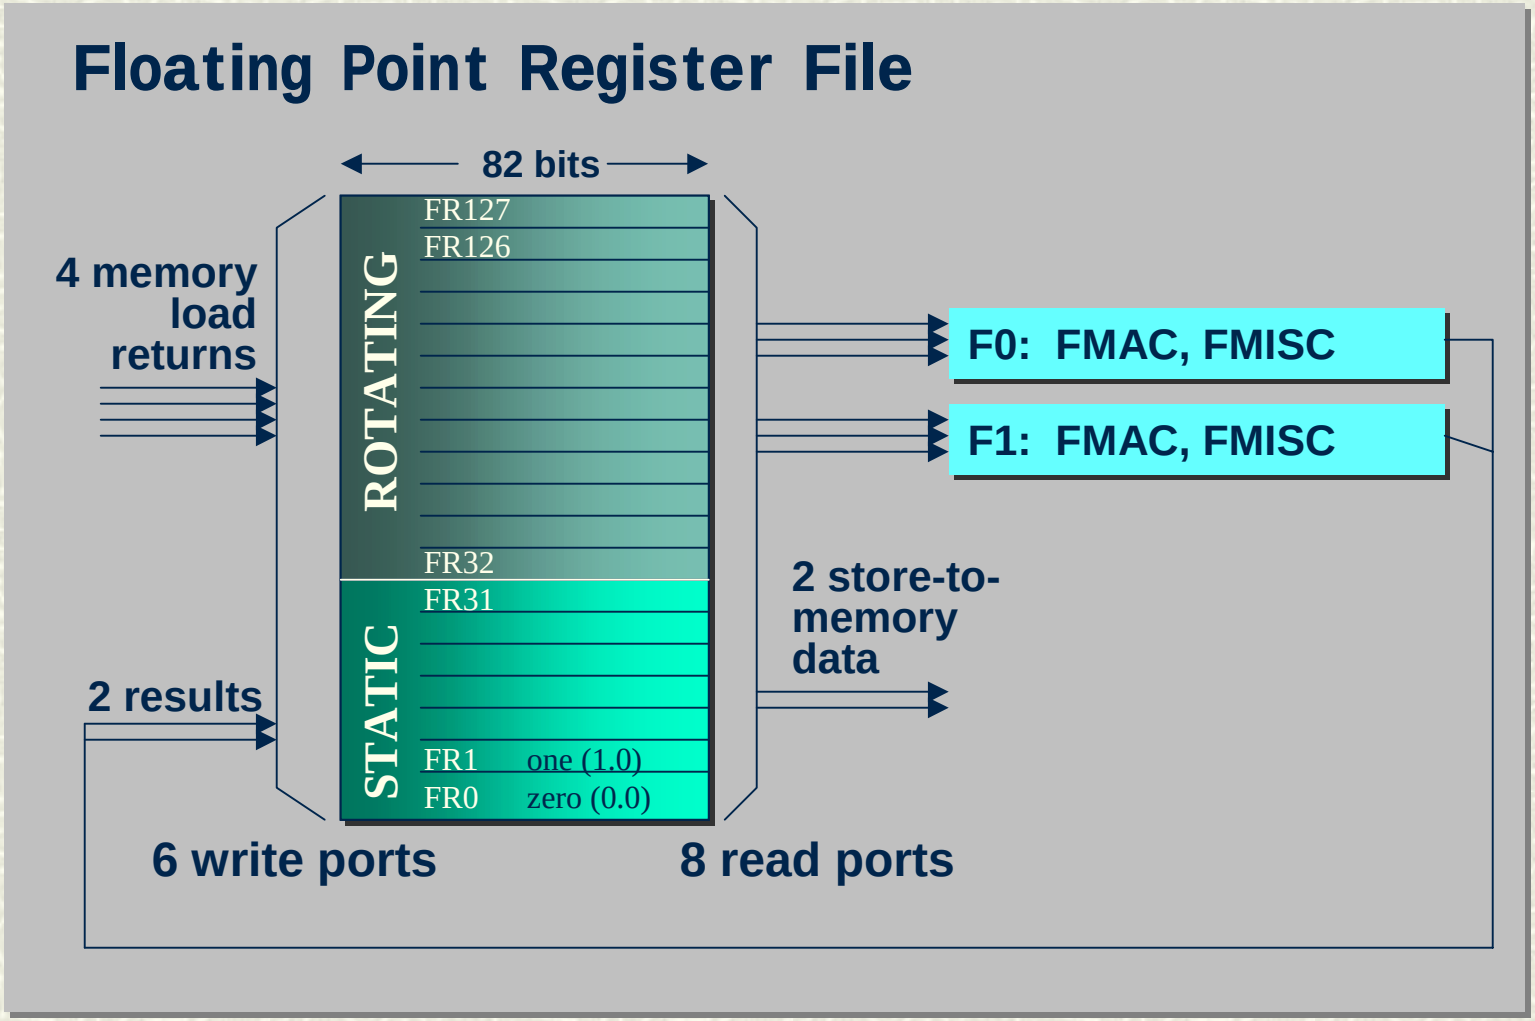
Floating Point Register File
82 bits
FR127
FR126
ROTATING
FR32
FR31
STATIC
FR1 one (1.0)
FR0 zero (0.0)
4 memory load returns
2 results
6 write ports
F0: FMAC, FMISC
F1: FMAC, FMISC
2 store-to-memory data
8 read ports

Itanium® 2 Processor Microarchitecture

Execution

Delays/Bypass

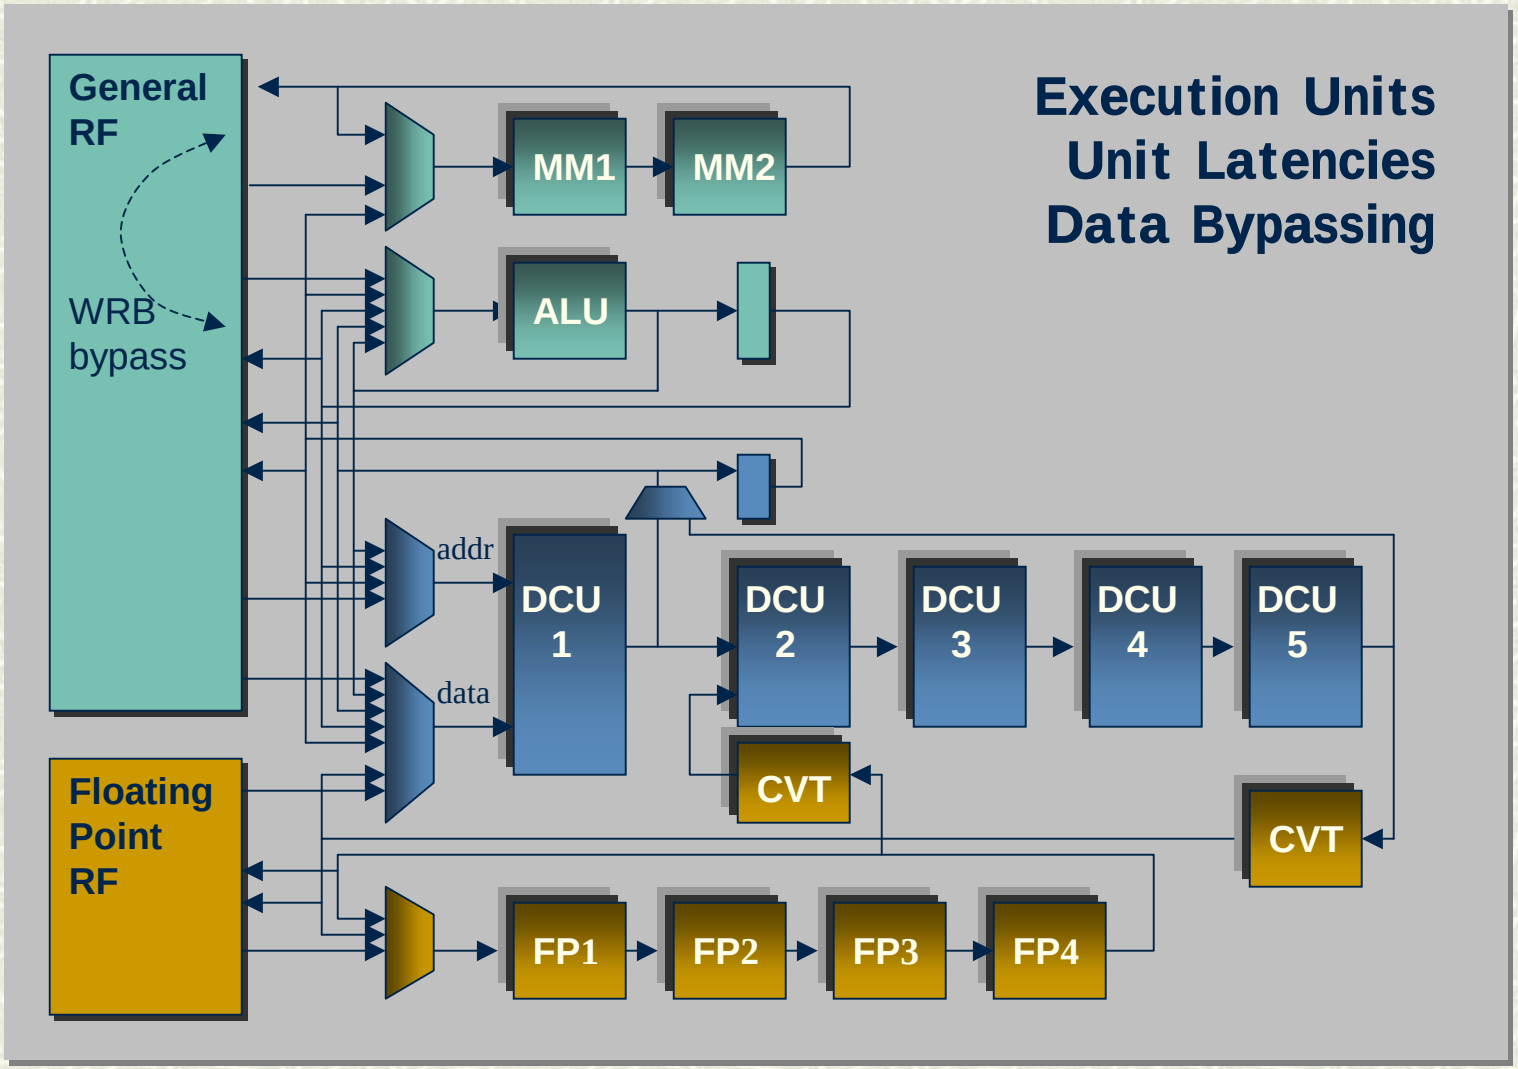

Stalls

Major Pipeline Stalls/Flushes

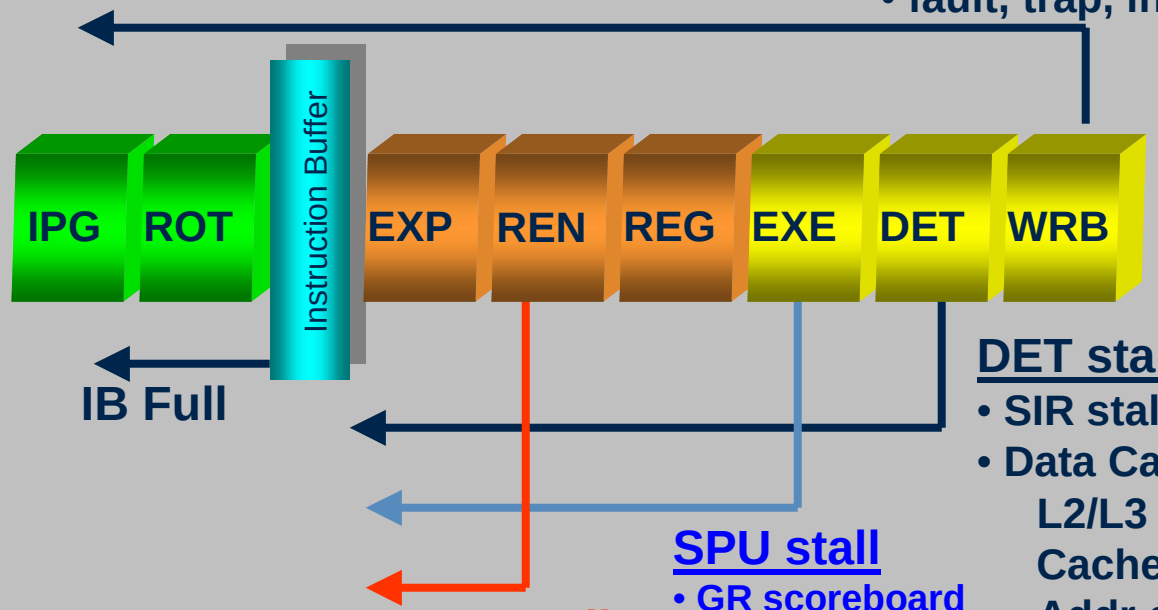

Itanium® 2 Processor Microarchitecture

3 Level Cache Organization

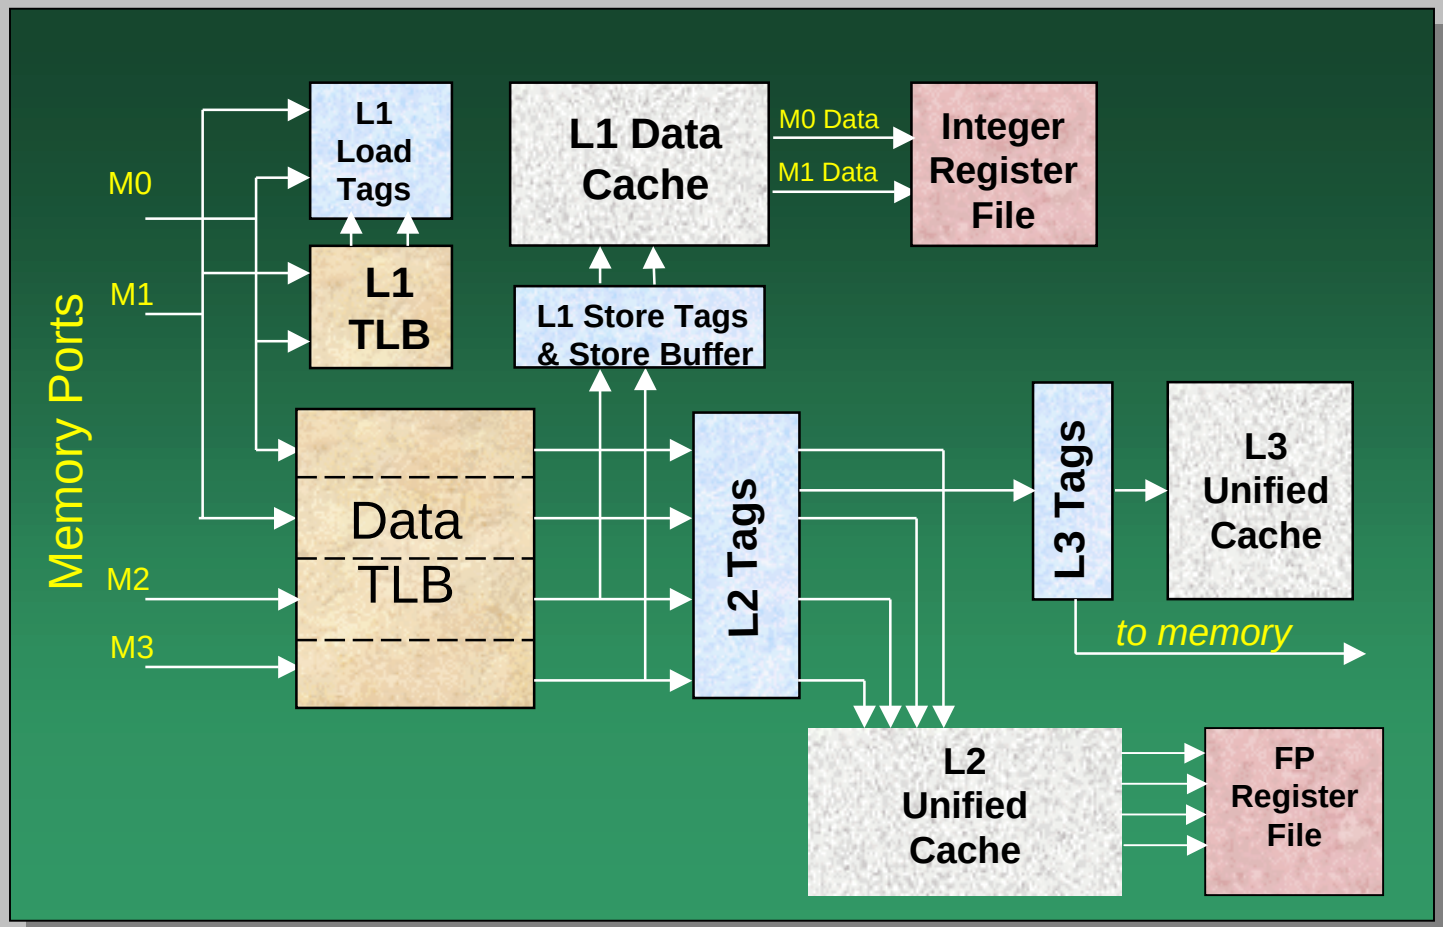

Itanium® 2 Processor Microarchitecture

Caches

L1d Cache – Integer Loads, Low Latency

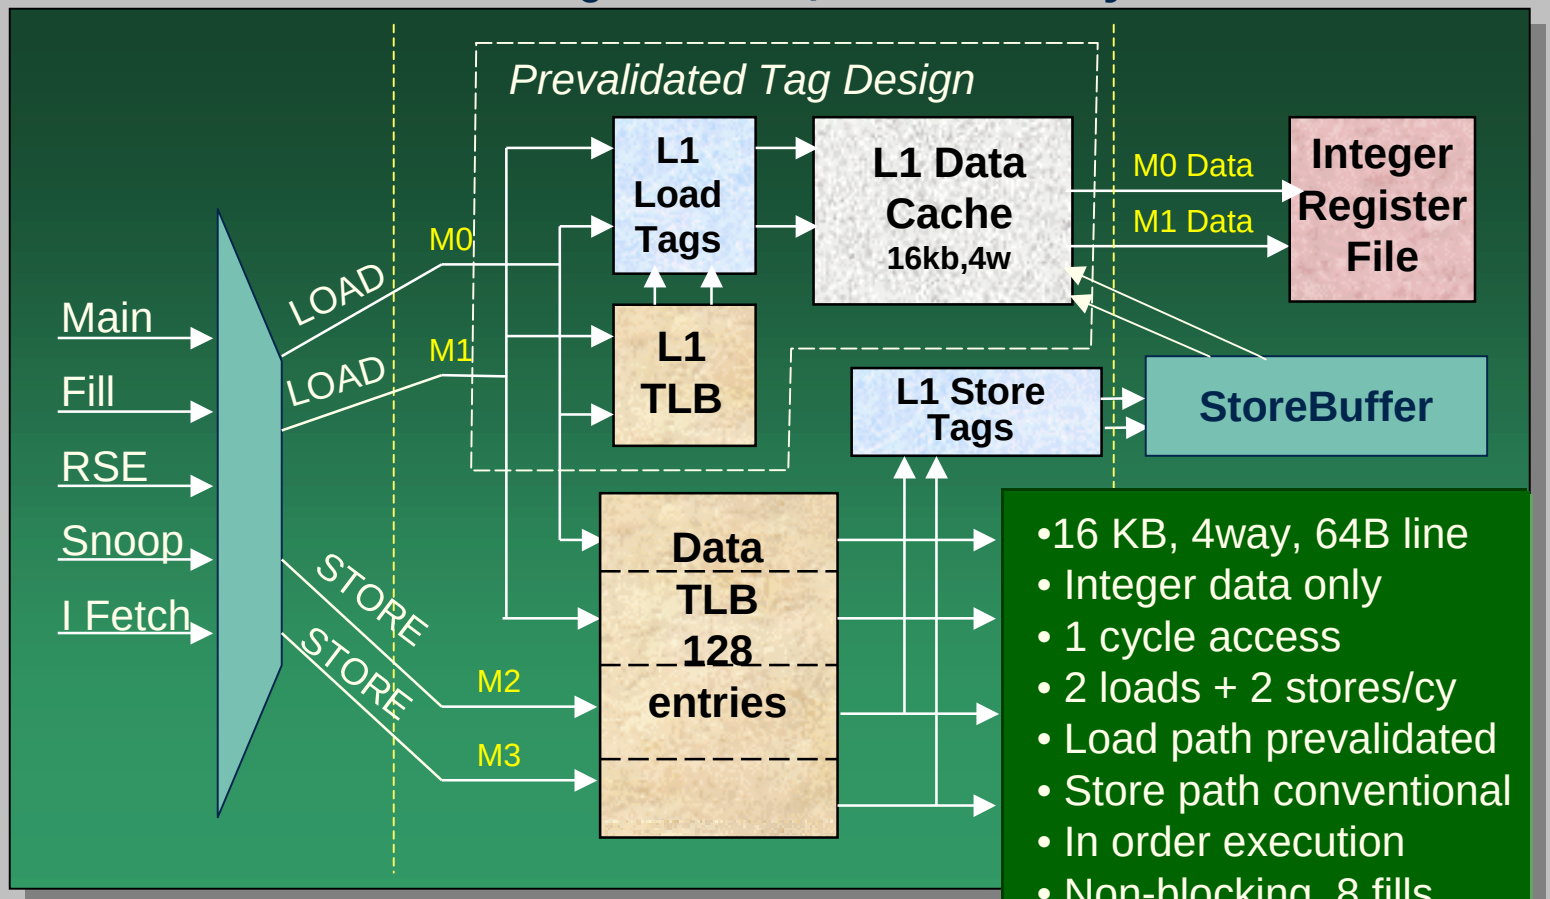

Prevalidated Cache Tag Design

Purpose: to provide a fast way hit

Conventional Tag

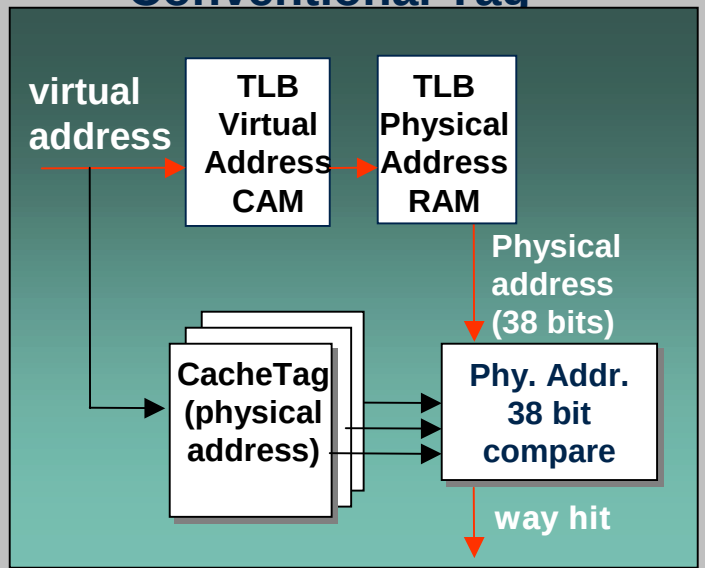

Logic Path: VA CAM
PA RAM
38 bit compare

Prevalidated Tag

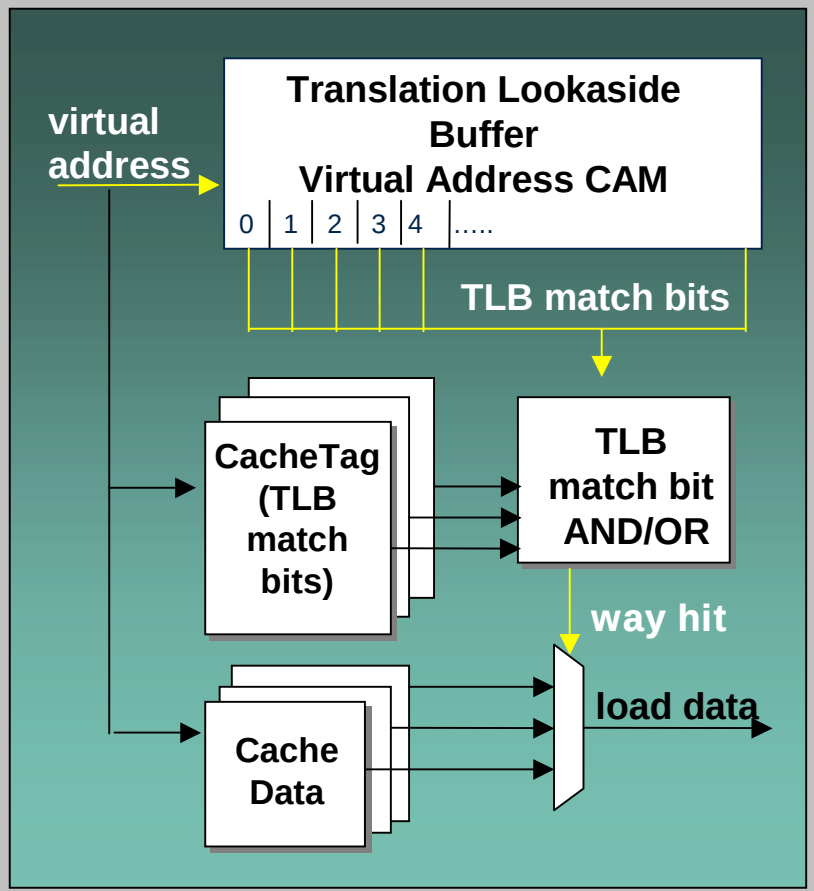

Logic Path: VA CAM
32 bit AND/OR

Advanced Load Address Table (ALAT)

Purpose: to allow loads to be executed speculatively ahead of possibly conflicting store operations

Implementation

- 32 fully associative entries
- max 20 bit physical address compare
- 2 loads/checks + 2 stores per cycle

How it Works

- On advanced loads -> allocate a new entry in ALAT
- On stores -> do a physical addr CAM and invalidate
- On checks -> read ALAT for a valid entry

On an Check ALAT Miss

- redo the load and resteer to the next instruction (ld.c)
- Or take a branch to recovery code (chk.a)

Itanium® 2 Processor Microarchitecture

Caches

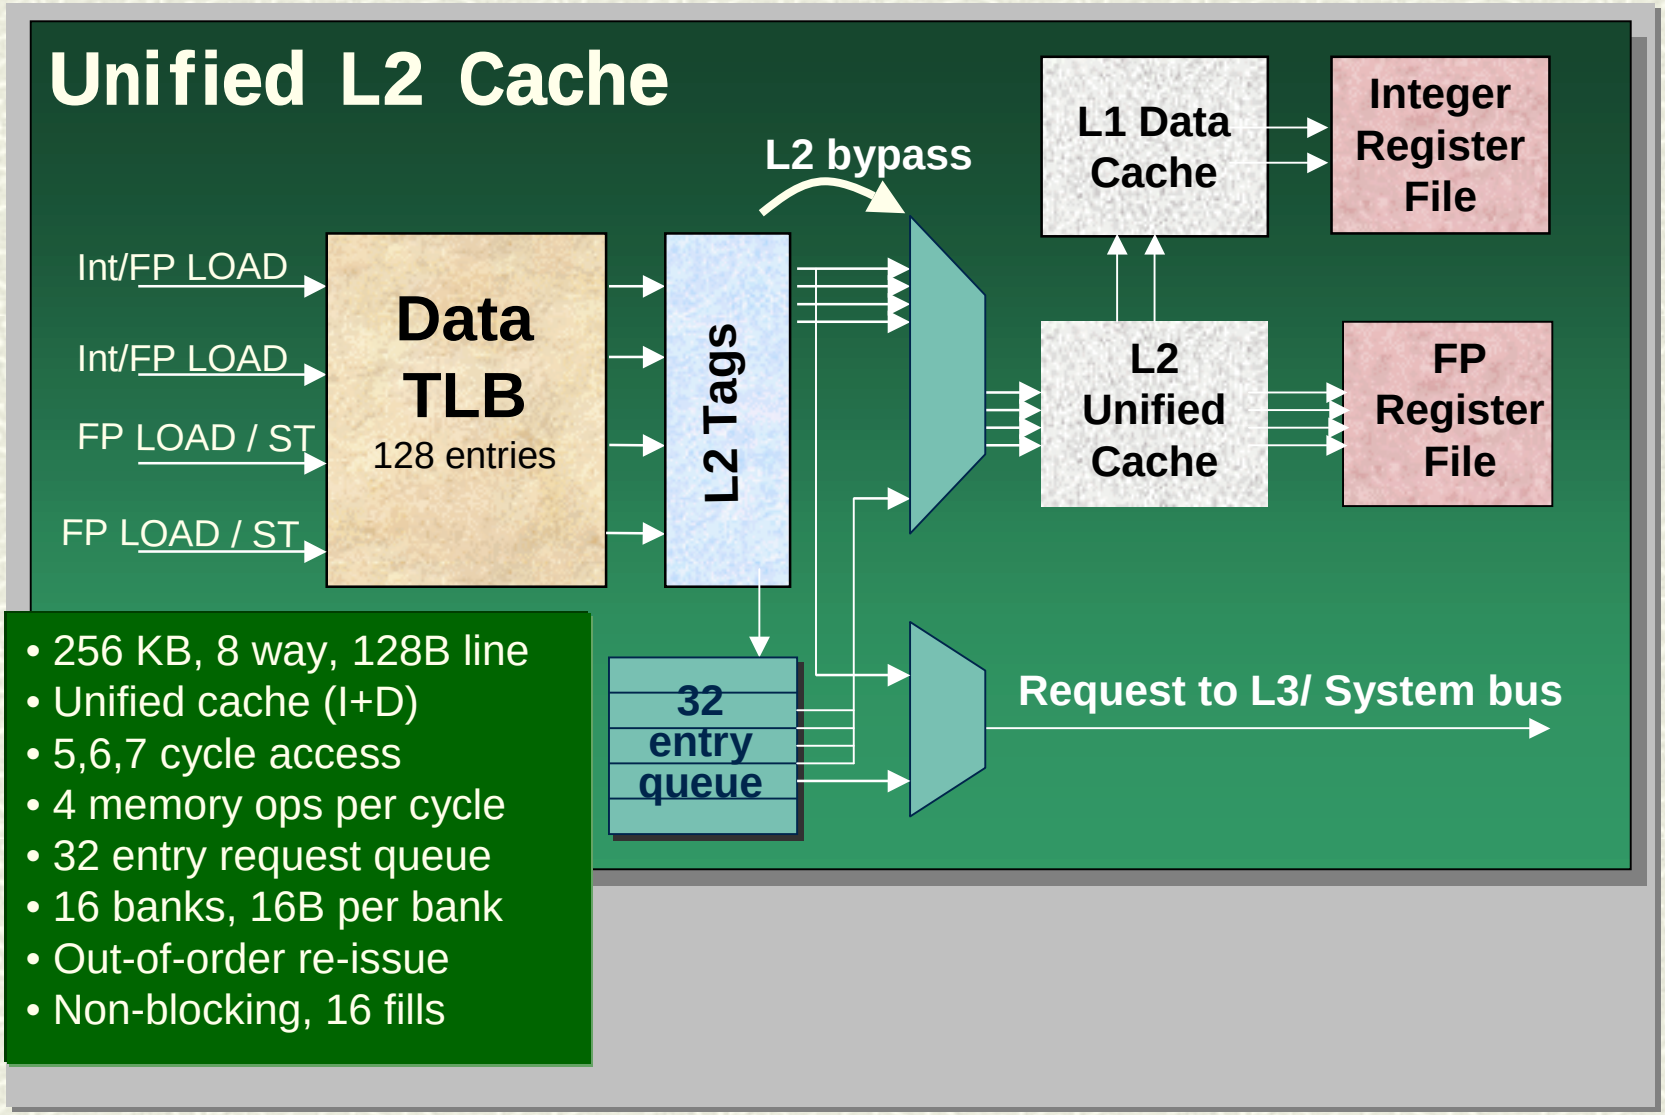

L2 Cache – The Floating-Point Streaming Engine!

STREAMING 4 FP OPERATIONS + 4 LOADS + 2 STORES

streaming at a 1 cycle rate with hits in the L2 cache

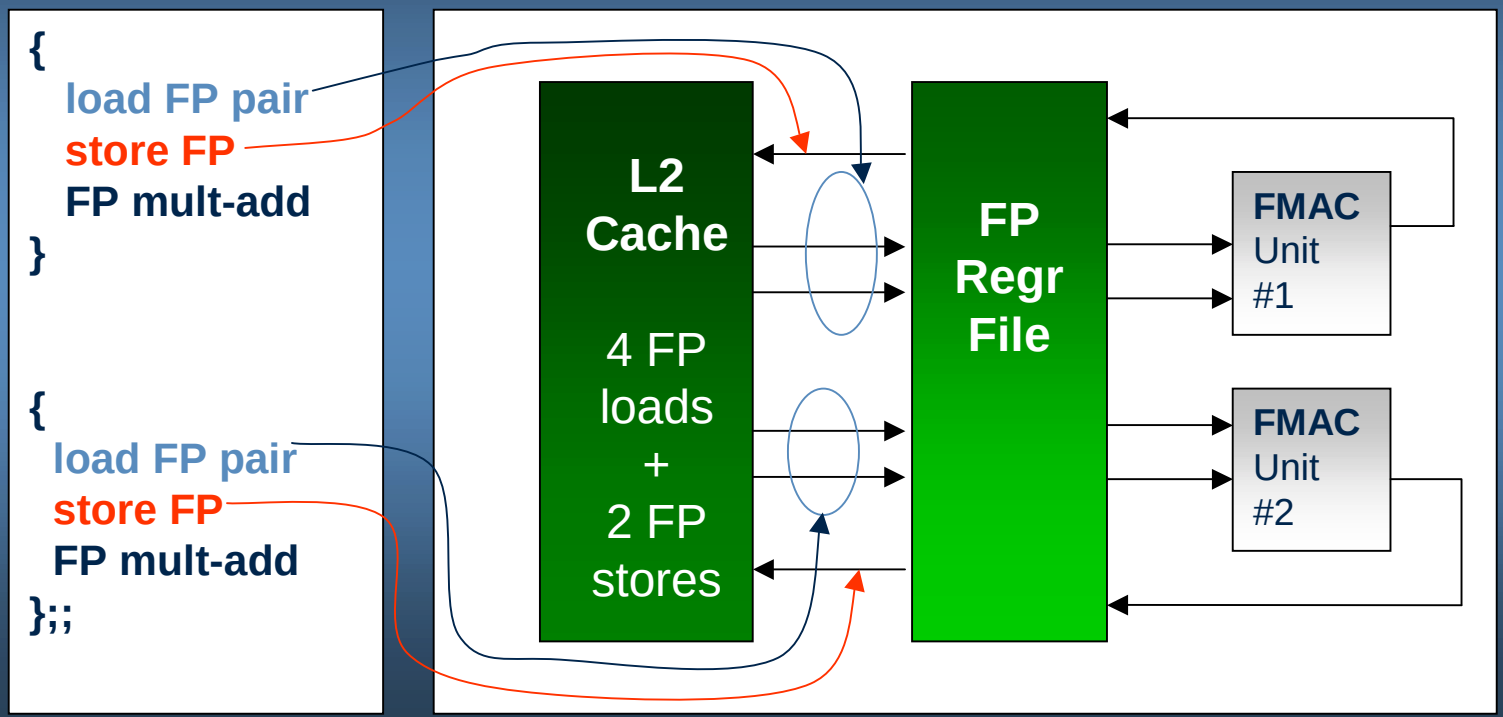

Itanium® 2 Processor Microarchitecture

Caches

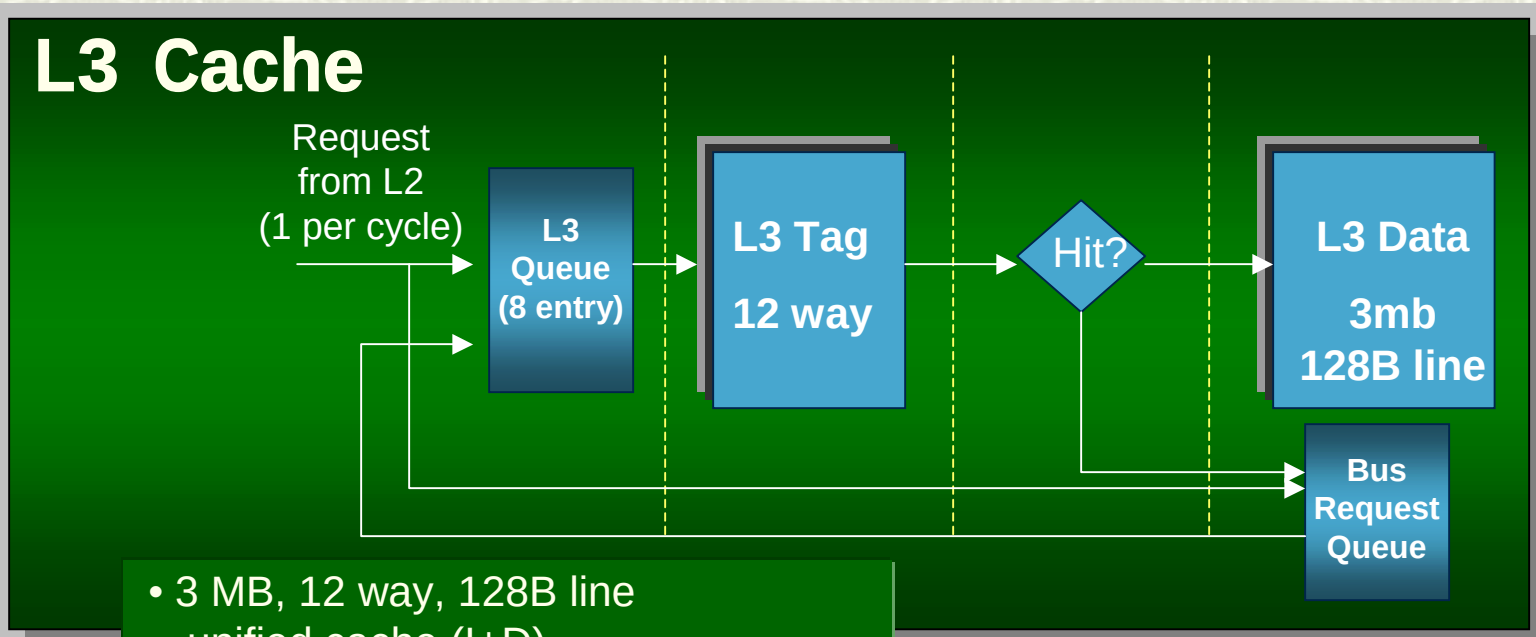

- 3 MB, 12 way, 128B line
- unified cache (I+D)
- On chip cache
- 12 cycle access (integer)
- Tag – 1 access per cycle
- Load data – 1 line per 4 cycles
- Store data – 1 line per 4 cycles
- Out-of-order re-issue
- Non-blocking, 16 entry miss buffer

Itanium® 2 Processor Microarchitecture

Caches

Major Data Paths

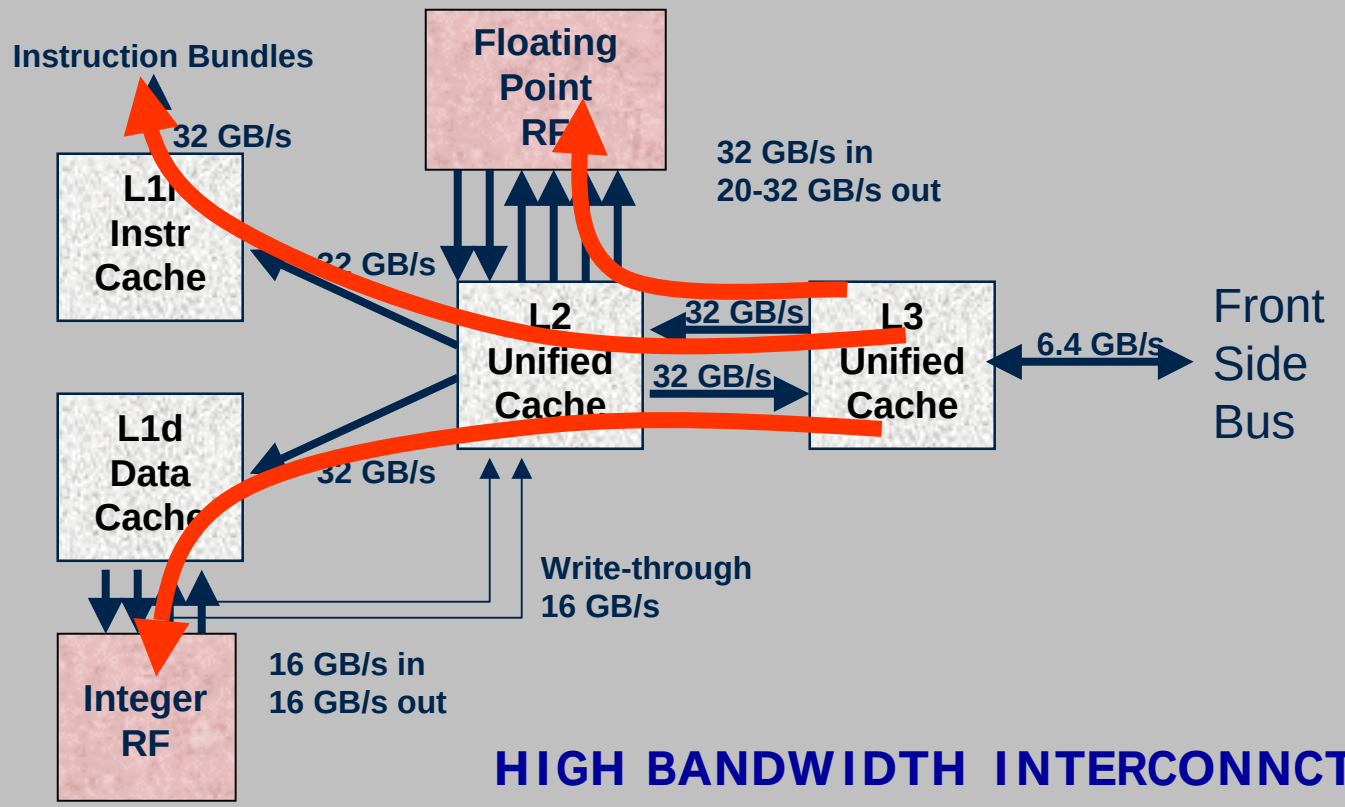

Memory Locality Hints

Purpose: to prevent cache pollution for non-temporal data

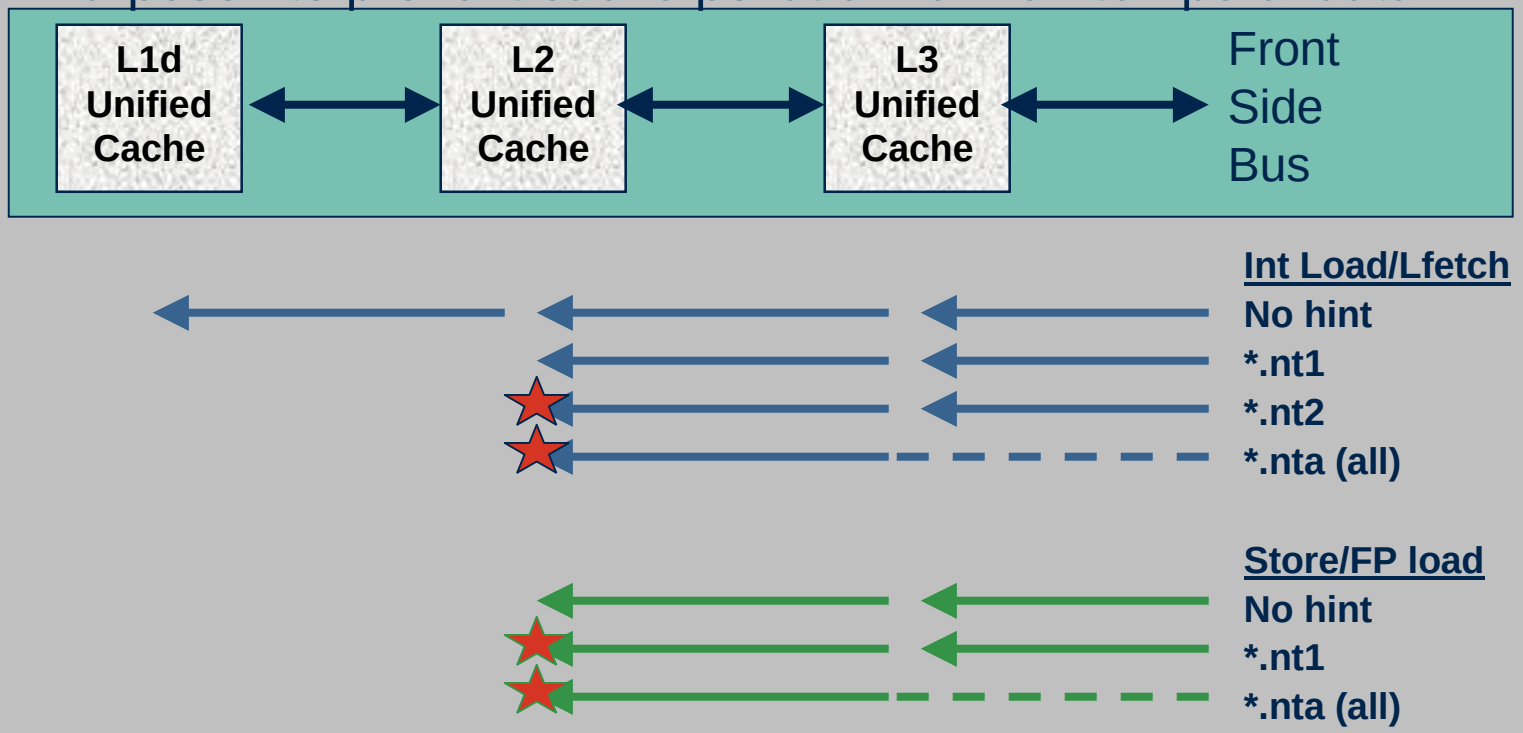

★ Replacement algorithm not updated

Itanium® 2 Processor Microarchitecture

Front Side Bus (FSB) Interface

- **Split bus** – independent address and data buses
- **Physical addressing – 50 bits**
- **Data bus**
  - 128 bits wide (16 Bytes)
  - 400 MT/s (double pumped 200MHz bus)
- **Queues**
  - supports 19 outstanding bus transactions
- **Multiprocessing Support**
  - glueless 4 CPU MP on one FSB
- **New bus functions**
  - “clean castout”, “read current”, “cleanse cache”

Itanium® 2 Processor Microarchitecture

Physical Design

Physical Design Stats

Frequency	1 GHz
Process	Intel 0.18 micron
	8” wafer, 6 layer Al metal
Die Size	421 mm^2 (full chip)
	142 mm^2 (processor core)
Number of Transistors	221 million
Power	130 watts @ 1GHz, 1.5v

Itanium® 2 Processor Microarchitecture

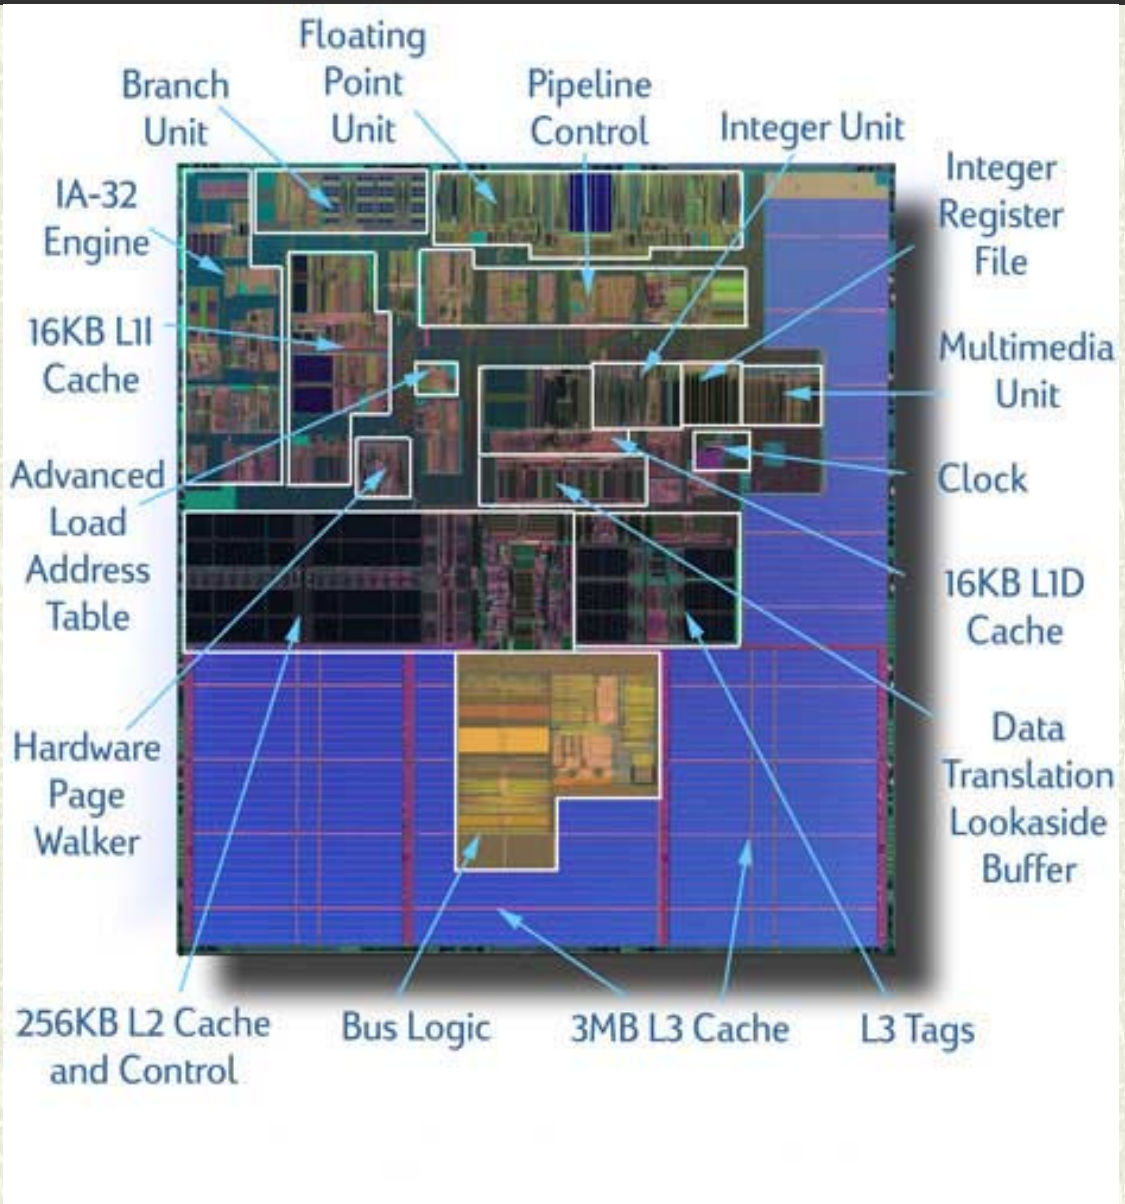

Physical Design

Itanium® 2 Processor Microarchitecture

University of Utah Math Dept Benchmark Suite

fluid dynamics and nonlinear PDEs codes

Run by HP with best compiled performance (no hand coding)

Power 4 is IBM p690 Regatta

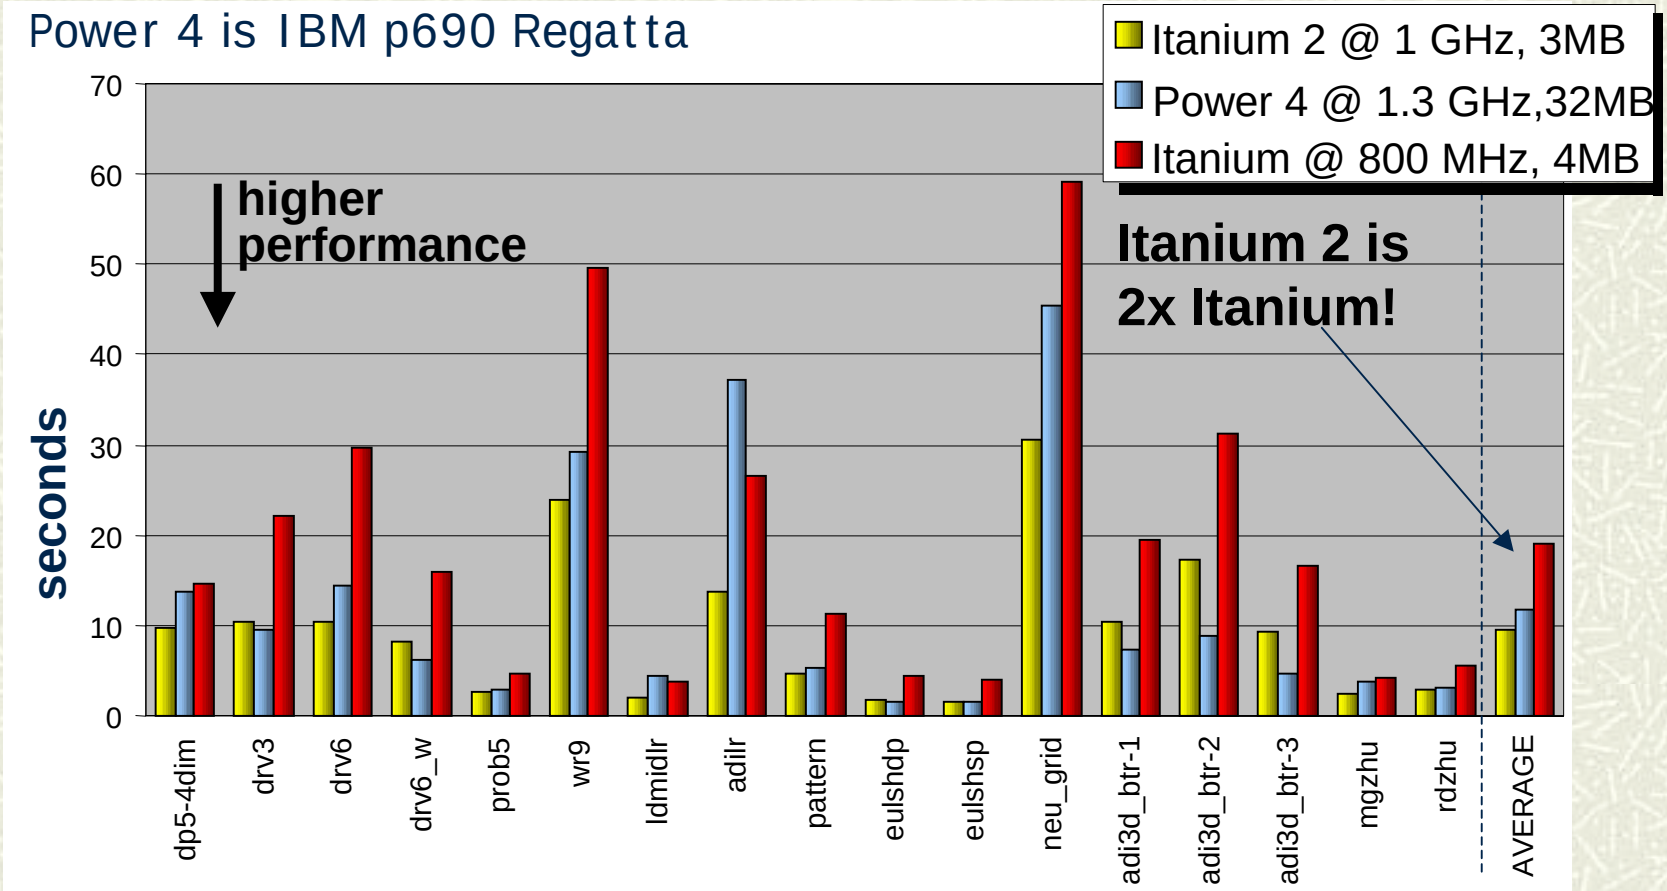

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, reference www.intel.com/procs/perf/limits.htm or call (U.S.) 1-800-628-8686 or 1-916-356-3104

Itanium® 2 Processor Microarchitecture

Performance

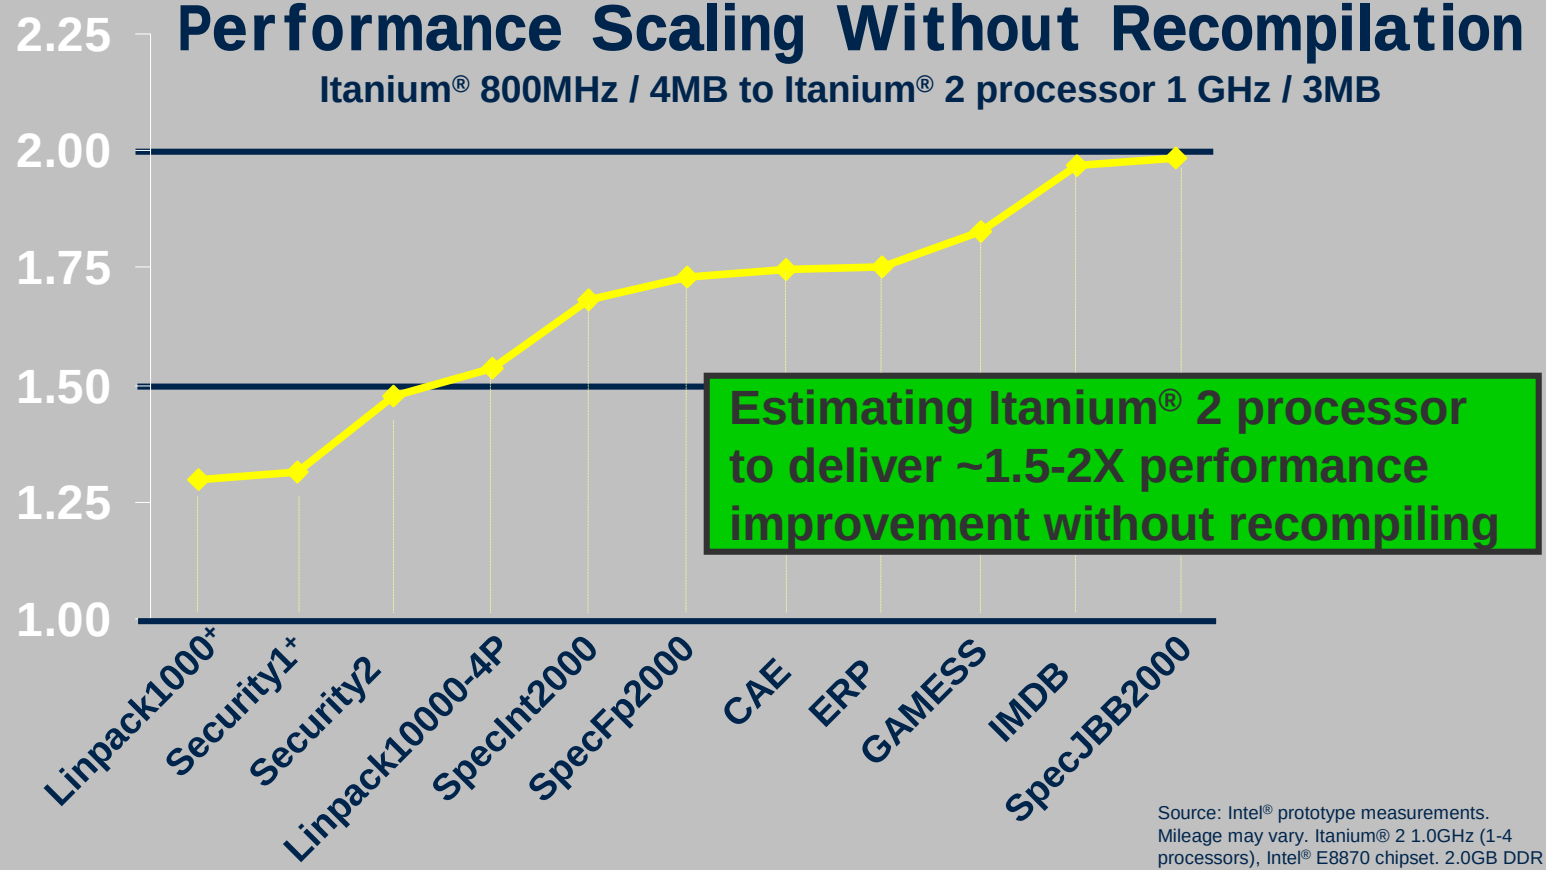

Source: Intel® prototype measurements. Mileage may vary. Itanium® 2 1.0GHz (1-4 processors), Intel® E8870 chipset. 2.0GB DDR SDRAM. Microsoft .net* Beta OS. Itanium® 800MHz (1-4 processors), Intel® 460GX chipset, 2.0GB SDRAM, Windows* 2000 64Bit Beta OS

+Linpack 1000, RSA already world class numbers on Itanium processor Data from measurements in Intel's labs, Itanium processor binaries on Itanium 2 processor

Itanium® 2 Processor Microarchitecture

Performance

Benchmark Performance

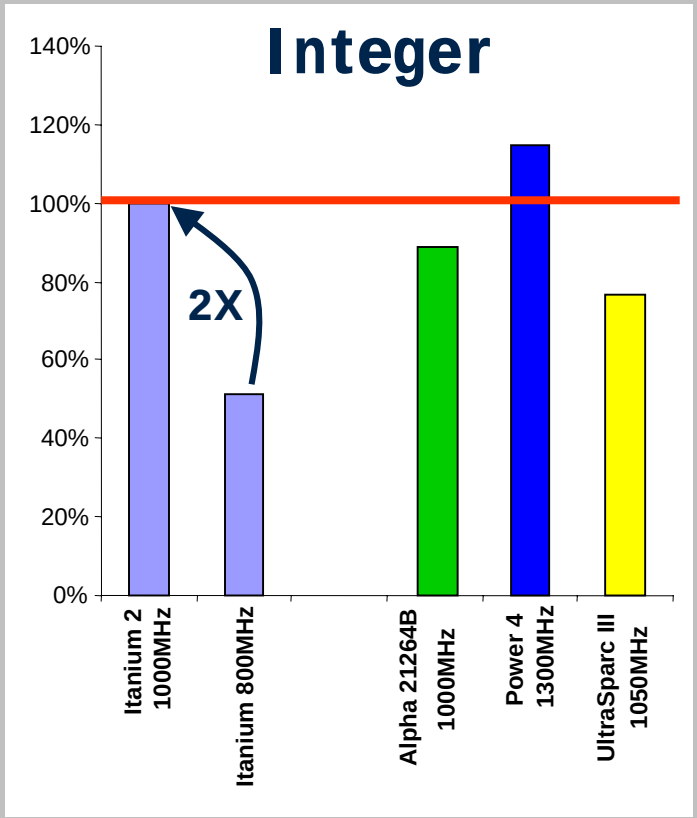

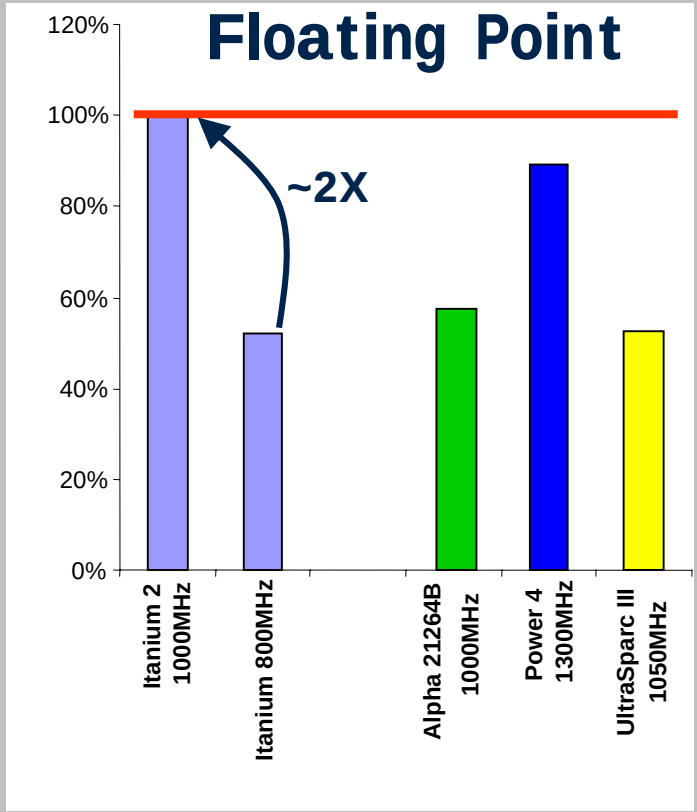

Competitive at 1 GHz **Leadership !**

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, reference www.intel.com/procs/perf/limits.htm or call (U.S.) 1-800-628-8686 or 1-916-356-3104

Itanium® 2 Processor Microarchitecture

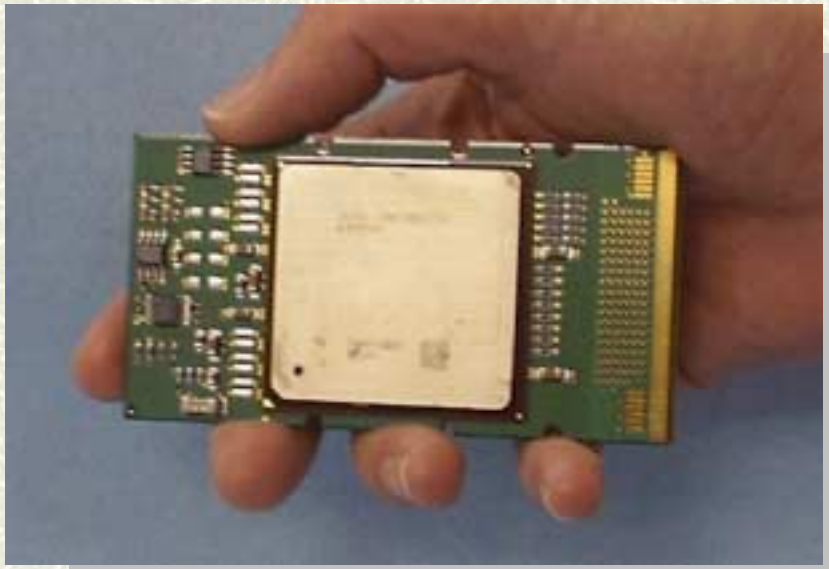

Further Information

- http://www.cpus.hp.com

  (papers on Itanium 2 and lots of pointers)

- http://developer.intel.com

  (resources on the Itanium Architecture)